

Physics-based Audio Synthesis in Computer Graphics

Adam Bognat, School of Computer Science

McGill University, Montreal

December, 2018

A thesis submitted to McGill University in partial fulfillment of the
requirements of the degree of

Master of Computer Science

©Adam Bognat, 2018

Abstract

Physics-based audio synthesis is finding increasingly diverse applications within computer graphics and animation. Conventional approaches based in modal synthesis are appropriate for simulations of rigid body contact sounds and large deformable objects, but are ill-suited for other primitives such as elastic rods and thin shells due to the non-linear effects that come with large displacements and the computational cost associated with highly refined meshes. In this thesis, we describe a software framework for the simulation of elastic rods and generation of sound from their motion that overcomes these limitations. A discrete elastic rod model based in Kirchhoff theory is used in simulating the rod dynamics, and a dipole source model is used to compute acoustic pressure values in the far-field. We consider a number of test systems to illustrate the fundamentals of the rod model and the variety of sounds that may be produced with a few free parameters.

Abréger

La synthèse audio basée sur la physique trouve des applications de plus en plus diverses dans les domaines de l’infographie et de l’animation. Les approches classiques basées sur la synthèse modale sont appropriées pour les simulations de sons de contact corps rigide et de grands objets déformables, mais sont mal adaptées à d’autres primitives telles que les tiges élastiques et les coques minces en raison des effets non linéaires liés aux grands déplacements et aux contraintes de calcul. coût associé à des mailles très raffinées. Dans cette thèse, nous décrivons un framework logiciel pour la simulation de tiges élastiques et la génération de son à partir de leur mouvement qui surmonte ces limitations. Un modèle de tige élastique discret basé sur la théorie de Kirchhoff simule la dynamique de la tige, et un modèle source dipolaire est utilisé pour calculer les valeurs de pression acoustique dans le champ lointain. Nous examinons un certain nombre de systèmes de test pour illustrer les principes fondamentaux du modèle à barres et la variété de sons pouvant être produits avec quelques paramètres libres.

Acknowledgements

I'd like to thank my supervisor Paul Kry, for his patient guidance throughout the course of my degree. Many meetings and coding sessions went into building the software framework that is the topic of this thesis. In particular, the plots in Chapter 3 of this thesis were produced using scripts he wrote.

I'd also like to thank Prakash Panangaden for his advice and assistance in switching fields from physics to computer science. He showed me that I did not have to be stuck, and that change was possible. Further thanks are owed to my mentor Nikunj Ragahuvanshi; working with him has given me a new perspective on what's possible for my career.

Finally, I'd like to thank my dear friend Guilherme Franzmann, for being a light in the dark and pillar of support for as long as I've known him.

Contents

1	Introduction	7
1.1	Audio Synthesis in Computer Graphics	8
1.2	Simulating Thin Rods	10
2	Discrete Elastic Rods	12
2.1	Smooth Adapted Framed Curves	12
2.2	Parallel-Transport	14
2.3	Discrete Adapted Framed Curves	14
2.4	Discrete Curvature	16
2.5	Parametrizing the Material Frames	17
2.6	Elastic Rod Dynamics	19
2.7	Force Computation	19
2.8	Damping Model	21
2.9	Validating the Force Expressions	22
2.10	Simulation Loop	22
3	Audio Synthesis	24
3.1	Sound Radiation Model	25
3.2	Dipole Source Implementation Details	26
4	Results	29
4.1	Ribbons, Polymers, and Slinkies	29

4.2 Rubber Bands and Metal Strings	32
5 Conclusion	38
5.1 Future Work	39
A Supplementary Materials	40

List of Figures

2.1	Discrete adapted framed curve	15
2.2	Material directors and reference frames	17
3.1	Velocity potential solutions	27
3.2	Interpolated buffer write	28
4.1	Three framed curves	30
4.2	Bent polymer	31
4.3	Cable resting under gravity	31
4.4	Displacement initial conditions	33
4.5	Cosine lobe force profile	33
4.6	Rubber band spectrograms: variable tension	34
4.7	Rubber band spectrograms: variable density	35
4.8	Rubber band spectrograms: variable initial conditions	35
4.9	Rubber band spectrograms: variable forcing	36
4.10	Nonlinear effects in metal strings	37

Chapter 1

Introduction

Procedural generation of sound finds many applications in computer music, interactive simulation, mixed reality applications, video games and computer animation. There are myriad approaches to sound synthesis, with the field lying in a rich intersection of physics, geometry, signal processing, and computer science. This thesis considers the physics-based simulation of sound from thin elastic rods, with which a number of non-trivial physical systems may be modelled. The work herein is primarily based on the framework described by Schweickart et al. [19].

We begin with an overview of sound rendering and audio synthesis in the context of computer graphics, discussing several approaches to generating sounds in pre-rendered and interactive audio-visual applications. We then narrow our focus to the topic of this thesis, giving an overview of the simulation of thin rods and methods of generating sounds from them. In particular, we address the limitations of common approaches and how the implementation considered in the thesis overcomes them.

The rest of the thesis is organized as follows: Chapter 2 introduces the physics of discrete elastic rods and the numerical integration of such systems; Chapter 3 introduces the dipole-based audio synthesis model used for sound generation and the pipeline used for offline sound rendering; Chapter 4 applies these concepts to the simulation of a number of physically-realistic systems; finally, Chapter 5 summarizes the work presented in

this thesis and gives avenues for future improvements. The appendix summarizes the supplemental materials, in the form of audio and video files, included with this thesis.

1.1 Audio Synthesis in Computer Graphics

Procedurally generated sound enriches computer graphics applications in ways that mere playback of pre-recorded sound does not. By synthesizing audio based on input from the simulation, the realism and fluidity of an animation, game or interactive mixed reality application is enhanced. Takala and Hahn’s introduction of the *sound rendering pipeline* [22] allowed for the procedural generation of audio events and their synchronization with images in the context of computer graphics and animation. In this framework, sound events could be instantiated (either procedurally in real-time or from pre-computed data), attached to objects in the scene, and finally rendered for auralization. In this way, the *modelling* of the sound could be separated from its playback, and allowed for additional integration of *modulation* of the sound to represent effects such as reverberation or spatialization.

Although there are a number of ways to generate sound, including traditional methods of synthesis in computer music and modulation of a pre-recorded base sound, we focus on *physics-based* methods of sound synthesis, wherein physically *motivated* (though not necessarily physically *accurate*) models of material vibration and wave propagation are used to synthesize sounds. A full solution involves solving the acoustic wave equation subject to appropriate boundary conditions set by the scene and motion of objects within it, which is possible in principle but prohibitively expensive and unnecessary for most applications in graphics and animation. Indeed, the goal of physics-based approaches is to construct models with a few free parameters and reduced degrees of freedom (relative to the complete phase space) that capture the salient features of the system of interest. Perceptual accuracy and computational efficiency take precedence over physical accuracy.

In a physics-based audio synthesis framework, there is usually a pre-computation phase that generates static data to be used in a dynamic synthesis phase, wherein input from the physical simulation drives a synthesis algorithm (though this synthesis phase is not necessarily capable of producing sounds in real-time). Examples of pre-computed data are the linear modes of vibration of a rigid body [17], basis functions of a multipole radiation source of a given order [8], and room impulse responses for the modelling of reverberation [18].

A wide variety of sounds may be accommodated by such a framework. While much consideration has been given to synthesizing rigid body contact sounds [15,16], including structures such as thin rods [19], shells [3], and arbitrarily complex vibration sources [8], there are physics-based models for the synthesis of sounds from fluids [13,24], and fire [4], as well as sounds due to fracturing [25], tearing [14], and crumpling [5].

Modal synthesis is a widely used approach to physics-based audio synthesis, that can more generally be applied to simulating the physics of deformable bodies [23]. For small amplitude vibrations about some equilibrium position, a deformable object behaves like a collection of decoupled damped harmonic oscillators, each with their characteristic frequency of oscillation and damping coefficient. By pre-computing these normal modes, one can efficiently synthesize the motion that results from an applied force, allowing for the real-time synthesis of audio and animation for hundreds of modes [9].

While readily applicable to extended near-rigid bodies, modal synthesis does not lend itself to thin, highly-deformable structures such as rods and shells. Large deformations change lead to time-varying normal modes, precluding precomputation, and capturing the modes of thin structures requires highly refined meshes, resulting in prohibitively expensive computation. Moreover, nonlinear effects, such as dispersion and pitch-shifting, are caused by large amplitude vibrations and cannot be modelled at all using linear modal synthesis.

To overcome these limitations, we consider a time-domain simulation of thin rods based in a Kirchhoff rod model. The state of the rod is represented by a limited number

of degrees of freedom, and no costly mesh simulation is required. For the purposes of sound generation, the rod is modelled as a collection of dipole sources, which leads to a straightforward acoustic pressure computation, assuming that the listener is in the far field. Our framework is implemented in Java; we made this choice because of the number of libraries readily available, including JOGL for graphics, Swing for UI, Vecmath for linear algebra computations, and a custom `GLEventListener` with an easily controllable trackball camera.

1.2 Simulating Thin Rods

In the context of computer graphics, thin rod models are often applied to simulating ropes, hair, and cloth, wherein the accurate and efficient resolution of hundreds of contacts is required [11]. The two standard rod models are the Kirchhoff and Cosserat models, which generalize the linear beam theory of Euler and Timoshenko to allow for geometric nonlinearities [12]. These models parametrize a rod state in terms of a discretized centerline composed of vertices and edges, with a material frame attached to each edge representing twisting and shearing about the centerline. The Kirchhoff model is a restriction of the Cosserat model to shear-free dynamics of the centerline.

Cosserat models have been considered in simulating hundreds of interacting strands [20, 21]. These models emphasize visually pleasing animations and thus are not well-suited to synthesis of realistic sounds. Bergou et al. [1, 2] have combined ideas from the field of discrete differential geometry to the theory of Kirchhoff rods to develop their own model of discrete elastic rods and viscous threads. Schweickert et al. [19] use this model as a starting point for physics-based synthesis of thin rod sounds, integrating the rod dynamics with a sophisticated contact model that allows for the generation of sound from hundreds of self-collisions for complex, extended systems. The goal of this thesis is to implement a simplified version of the framework developed in Schweickert et al. [19], demonstrating the fundamentals of Bergou’s discrete elastic rod model and how it may

be combined with a simple sound radiation model to produce a variety of sounds in a reasonably controlled way.

Chapter 2

Discrete Elastic Rods

In this chapter, we define the physical model underlying the description of a discrete elastic rod. We define the kinematics of an adapted framed curve, discuss the generation of frames by means of parallel transport, and define the dynamics of the curve by means of a linear elasticity model. We give an overview of the high-level implementation of the relevant force computation in our simulation framework, concluding with a discussion of some systematic ways of validating our force expressions by means of finite difference approximations. We follow the discussion presented in Bergou et al. [1, 2] and Jawed et al. [10].

2.1 Smooth Adapted Framed Curves

We start with the smooth setting to build some intuition about the underlying geometric concepts that entail a description of elastic rods. A *smooth adapted framed curve* is defined by a time-dependent centerline and set of material frames

$$\mathbf{x}(s, t) \in \mathbb{R}^3, \quad [\mathbf{d}_1(s, t) \ \mathbf{d}_2(s, t) \ \mathbf{d}_3(s, t)] \in \text{SO}(3),$$

where s is a material coordinate parametrizing distance along the centerline and t is the time. That the frame is *adapted* to the centerline means that one of the vectors comprising

the orthonormal basis is parallel to the tangent $\mathbf{t} = \mathbf{x}'/|\mathbf{x}|$ of the curve, where $'$ denotes differentiation with respect to s . Physically, this choice amounts to disallowing shearing of the material frame. Conventionally, we take $\mathbf{d}_3 = \mathbf{t}$; then $\mathbf{d}_1, \mathbf{d}_2$ span a plane normal to the centerline at each point on the curve.

Denoting time differentiation with an overdot, the kinematics of the centerline are determined by its velocity $\dot{\mathbf{x}}(s, t)$ and deformation gradient $\mathbf{x}'(s, t)$. The deformation gradient, in turn, defines the relative axial strain $\epsilon = |\mathbf{x}'|/|\bar{\mathbf{x}}'| - 1$, where quantities with an overbar are associated with the undeformed equilibrium configuration of the centerline.

The kinematics of the material frame are defined by infinitesimal rotation, namely,

$$\dot{\mathbf{d}}_\alpha = \boldsymbol{\omega} \times \mathbf{d}_\alpha, \quad \alpha = 1, 2, 3,$$

where the angular velocity $\boldsymbol{\omega}$ is given by

$$\boldsymbol{\omega} = \mathbf{t} \times \dot{\mathbf{t}} + \omega_3 \mathbf{t}.$$

Geometrically, $\mathbf{t} \times \dot{\mathbf{t}}$ is the component of rotation needed to keep the frame adapted, with the remaining component inducing a rotation about the tangent with rotation rate $\omega_3 = \dot{\mathbf{d}}_1 \cdot \mathbf{d}_2$.

The spatial evolution of the material frames satisfies an analogous equation, given by

$$\mathbf{d}'_\alpha = \boldsymbol{\Omega} \times \mathbf{d}_\alpha, \quad \alpha = 1, 2, 3,$$

where the Darboux vector $\boldsymbol{\Omega}$ is given by

$$\boldsymbol{\Omega} = \mathbf{t} \times \mathbf{t}' + m\mathbf{t}.$$

In this decomposition, rotation about the curvature binormal $\kappa \mathbf{b} = \mathbf{t} \times \mathbf{t}'$ keeps the frame adapted along the centerline, while rotation about the tangent by the twist $m = \mathbf{d}'_1 \cdot \mathbf{d}_2$ accommodates the material twist in the rod.

We can define the smooth material curvatures (κ_1, κ_2) , to be used as a measure of the bending strain of the rod, as follows: since $\{\mathbf{t}, \mathbf{d}_1, \mathbf{d}_2\}$ form an orthonormal frame, and $\mathbf{t}' \perp \mathbf{t}$, we have that $\mathbf{t}' \in \text{span}\{\mathbf{d}_1, \mathbf{d}_2\}$. Thus we set

$$\kappa_1 = \mathbf{t}' \cdot \mathbf{d}_1, \quad \kappa_2 = \mathbf{t}' \cdot \mathbf{d}_2.$$

Using elementary vector product identities, we find

$$\kappa_1 = (\kappa \mathbf{b}) \cdot \mathbf{d}_2, \quad \kappa_2 = -(\kappa \mathbf{b}) \cdot \mathbf{d}_1.$$

These expressions will motivate the definition of the discrete material curvatures later on.

2.2 Parallel-Transport

To describe the kinematics in the discrete case, we need the notion of parallel transport. The parallel transport $P_{\mathbf{r}_1}^{\mathbf{r}_2}$ from unit vector \mathbf{r}_1 to unit vector \mathbf{r}_2 is given by

$$P_{\mathbf{r}_1}^{\mathbf{r}_2} = R(\mathbf{r}_1 \times \mathbf{r}_2 / |\mathbf{r}_1 \times \mathbf{r}_2|, \angle(\mathbf{r}_1, \mathbf{r}_2)),$$

where $R(\mathbf{r}, \theta)$ is the rotation about vector \mathbf{r} by angle θ . Thus $P_{\mathbf{r}_1}^{\mathbf{r}_2}$ is the *minimal* rotation that takes \mathbf{r}_1 into \mathbf{r}_2 .

2.3 Discrete Adapted Framed Curves

Formally, a discrete adapted framed curve consists of a centerline defined by $n + 2$ vertices $\mathbf{x}_i \in \mathbb{R}^3$ and $n + 1$ edges $\mathbf{e}^i = \mathbf{x}_{i+1} - \mathbf{x}_i$ (Figure 2.1). To each edge is associated an




Figure 2.1: Notation used in the definition of discrete adapted framed curves.

orthonormal frame $[\mathbf{d}_1^i \ \mathbf{d}_2^i \ \mathbf{d}_3^i] \in \text{SO}(3)$ such that $\mathbf{d}_3^i = \mathbf{t}^i = \mathbf{e}^i/|\mathbf{e}^i|$. The kinematic state is defined by the vertex positions and velocities, and the relative axial-strain of the smooth case is replaced by an edge-based relative axial strain, given by $\epsilon^i = |\mathbf{e}^i|/|\bar{\mathbf{e}}^i|$, where we remind the reader that overbar quantities are associated with the undeformed equilibrium configuration of the centerline.

We can express the kinematics of the discrete frames by a composition of parallel-transport and a rotation about the tangent. Assuming we have reasonable definitions of the discrete angular velocity ω_3^i (to be discussed later), the i 'th frame at time t_k is obtained from the frame at time t_{k-1} according to

$$\mathbf{d}^i(t_k)_\alpha = \mathbf{R}(\mathbf{t}^i(t_k), \omega_3^i(t_k)) \circ \mathbf{P}_{\mathbf{t}^i(t_{k-1})}^{\mathbf{t}^i(t_k)} \mathbf{d}_\alpha^i(t_{k-1}).$$

In words, we parallel transport the old frame to the new tangent, and then rotate about the new tangent. The spatial evolution is governed analogously, except the parallel-transport takes us between adjacent edges at a *fixed* time, namely,

$$\mathbf{d}_\alpha^i = \mathbf{R}(\mathbf{t}^i(t_k), m_i(t_k)) \circ \mathbf{P}_{\mathbf{t}^{i-1}(t_k)}^{\mathbf{t}^i(t_k)} \mathbf{d}_\alpha^{i-1}(t_k).$$

Note that the discrete frames are updated using *finite* rotations. We henceforth refer to $\mathbf{P}_{\mathbf{t}^i(t_{k-1})}^{\mathbf{t}^i(t_k)}$ as the *time-parallel* transport operator and $\mathbf{P}_{\mathbf{t}^{i-1}(t_k)}^{\mathbf{t}^i(t_k)}$ as the *space-parallel* transport operator.

2.4 Discrete Curvature

In this section we briefly describe some notions of discrete curvature. We follow the discussion presented by Jawed et al. [10]. Consider 3 vertices $\mathbf{v}_{i-1}, \mathbf{v}_i, \mathbf{v}_{i+1}$ and their corresponding edges $\mathbf{e}^{i-1}, \mathbf{e}^i$. We seek a sensible definition of a discrete curvature binormal $\kappa \mathbf{b}_i$. To this end, define the turning angle ϕ_i between the edges by

$$\cos(\phi_i) = \mathbf{t}^{i-1} \cdot \mathbf{t}^i, \quad |\sin(\phi_i)| = |\mathbf{t}^{i-1} \times \mathbf{t}^i|.$$

Using this and the Voronoi length $l_i = \frac{1}{2}(|\mathbf{e}^{i-1}| + |\mathbf{e}^i|)$ associated with the i 'th vertex, we can define a discrete pointwise curvature $\hat{\kappa}_i$ at the central vertex by

$$\hat{\kappa}_i = \frac{2}{l_i} \tan(\phi_i/2).$$

Let R_i be the radius of the osculating circle tangent to the edges $\mathbf{e}^{i-1}, \mathbf{e}^i$. Then $\hat{\kappa}_i = 1/R_i$.

From the discrete pointwise curvature we can define a discrete integrated curvature $\kappa_i = \hat{\kappa}_i l_i$. We combine this with the definition of the discrete binormal vector,

$$\mathbf{b}_i = \frac{\mathbf{t}^{i-1} \times \mathbf{t}^i}{|\mathbf{t}^{i-1} \times \mathbf{t}^i|},$$

to obtain the discrete integrated curvature binormal at the central vertex, namely,

$$\kappa \mathbf{b}_i = \frac{2 \mathbf{t}^{i-1} \times \mathbf{t}^i}{1 + \mathbf{t}^{i-1} \cdot \mathbf{t}^i}.$$

We can now define the vertex-based curvatures of the material frame, which will be used to define the bending strains of the elastic rod in a subsequent section. Given material frames $\mathbf{d}_1^{i-1}, \mathbf{d}_2^{i-1}, \mathbf{d}_1^i, \mathbf{d}_2^i$, the set $\{\kappa \mathbf{b}_i, \frac{1}{2}(\mathbf{d}_1^{i-1} + \mathbf{d}_1^i), \frac{1}{2}(\mathbf{d}_2^{i-1} + \mathbf{d}_2^i)\}$ constitutes an orthonormal frame at the central vertex. With this, we may define the discrete material




Figure 2.2: Material directors and reference frames. The tangent \mathbf{d}_3 points out of the page.

curvatures (κ_1, κ_2) by

$$\begin{aligned}\kappa_1 &= \frac{1}{2}(\mathbf{d}_2^{i-1} + \mathbf{d}_2^i) \cdot \kappa \mathbf{b}, \\ \kappa_2 &= -\frac{1}{2}(\mathbf{d}_1^{i-1} + \mathbf{d}_1^i) \cdot \kappa \mathbf{b}.\end{aligned}$$

These expressions are the obvious discretization of the smooth material curvatures defined above.

2.5 Parametrizing the Material Frames

In this section we discuss a convenient parametrizing of the material frames by means of an appropriate choice of reference frames.

Consider an arbitrarily chosen orthonormal frame on the first edge of a discrete curve. By space-parallel transporting this initial frame along the curve, we get an initial set of reference frames. These frames will, in general, differ from the true material frames of the curve. The material frame at each edge will be related to the reference frame by a rotation about some edge angle θ^i (Figure 2.2). Denoting the reference frames quantities with an

underbar, we have

$$\begin{aligned}\mathbf{d}_1^i &= (\cos(\theta^i)\underline{\mathbf{d}}_1^i + (\sin(\theta^i))\underline{\mathbf{d}}_2^i, \\ \mathbf{d}_2^i &= -(\sin(\theta^i)\underline{\mathbf{d}}_1^i + (\cos(\theta^i))\underline{\mathbf{d}}_2^i.\end{aligned}$$

With this parametrization, the discrete tangential angular velocity ω_3 and twist m at time t_k are given by

$$\begin{aligned}\omega_3^i(t_k) &= \theta^i(t_k) - \theta^i(t_{k-1}) + \underline{\omega}_3^i(t_k), \\ m_i(t_k) &= \theta^i(t_k) - \theta^{i-1}(t_k) + \underline{m}_i(t_k).\end{aligned}$$

If we update the reference frames at each time-step by time-parallel transport, we get that $\underline{\omega}_3 = 0$. This choice makes update of the material frame straightforward. Having stepped the vertex positions and edge angles, we get the new material frames by time-parallel transporting the old reference frames and rotating the new reference frames about the tangent by the new edge angle. Expressed in terms of operators, we have

$$\mathbf{d}_\alpha^i(t_k) = \mathbf{R}(\mathbf{t}^i(t_k), \theta^i(t_k)) \circ \mathbf{P}_{\mathbf{t}(t_{k-1})}^{\mathbf{t}(t_k)} \circ \mathbf{R}(\mathbf{t}^i(t_{k-1}), -\theta^i(\mathbf{t}_{k-1})) \mathbf{d}_\alpha^i(t_{k-1}).$$

The use of time-parallel reference frames makes the computation of the twist m , used in computing the energy, straightforward. The material twist of the rod determines the extent to which the frames deviate from parallel. Thus, we get the twist by computing the angle between the material frames on edge i to the space-parallel transport of the material frame on edge $i - 1$ to edge i .

2.6 Elastic Rod Dynamics

The rod dynamics are governed by the equations of linear elasticity. The elastic stretching, twisting and bending energies are given by

$$\begin{aligned} E_s &= \frac{1}{2} \sum_{j=0}^n k_s^j (\epsilon^j)^2 |\bar{\mathbf{e}}^j|, \\ E_t &= \frac{1}{2} \sum_{j=1}^n \beta_j \frac{(m_j - \bar{m}_j)^2}{\bar{l}_j}, \\ E_b &= \frac{1}{2} \sum_{j=1}^n \frac{1}{\bar{l}_j} (\kappa_j - \bar{\kappa}_j)^T \mathbf{B}_j (\kappa_j - \bar{\kappa}_j), \end{aligned}$$

where $\bar{l}_j = \frac{1}{2}(|\bar{\mathbf{e}}^{j-1}| + |\bar{\mathbf{e}}^j|)$ denotes the Voronoi length associated to the j 'th vertex and the symmetric quadratic form \mathbf{B} measures the anisotropic bending response along the material axes of the curve. The subtraction of the \bar{m} and $\bar{\kappa}$ allows for the accommodation of curved and twisted equilibrium configurations.

The force stiffnesses will generally depend on the geometry of the centerline's cross section. Assuming each edge has an elliptical cross-section with major and minor axes a^i , b^i , and area A^i , the stiffnesses can be written as

$$k_s^i = EA^i, \quad \beta_i = \frac{GA_i}{4}(a_i^2 + b_i^2), \quad B_i = \frac{E}{4A_i} \text{diag}(a_i^2, b_i^2),$$

where E and G denote the Young's modulus and shear modulus of the material, respectively. It is common to give the Young's modulus and Poisson ratio ν , from which we get $G = \frac{E}{2(1 + \nu)}$.

2.7 Force Computation

We obtain forces from energies by taking partial derivatives with respect to the vertex positions and edge angles. In our simulation framework, each force element is responsi-

ble for applying its contribution to the net force on its degrees of freedom. To illustrate, consider a minimal system consisting of 3 vertices and 2 edges which define 2 stretching force elements, 1 bending force element, and 1 twisting force element. The degrees of freedom associated to the bending element consist of the 3 vertex positions $\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2$ and 2 edge angles θ_0, θ_1 , from which one may construct a state vector $\mathbf{x} \in \mathbb{R}^{11}$ by

$$\mathbf{x} = (v_{0x}, v_{0y}, v_{0z}, \theta_0, v_{1x}, v_{1y}, v_{1z}, \theta_1, v_{2x}, v_{2y}, v_{2z}).$$

Then the contributions of these force elements to the net force on the corresponding degrees of freedom are given by

$$\begin{aligned}\mathbf{F}_b &= - \left(\frac{\partial E_b}{\partial \mathbf{x}} \right)^T, \\ \mathbf{F}_t &= - \left(\frac{\partial E_t}{\partial \mathbf{x}} \right)^T.\end{aligned}$$

Note that the bending and twisting force elements both apply forces to the vertices and moments to the edges. To get the stretching forces, one would consider the state vectors $\mathbf{x}_{s0} = (v_{0x}, v_{0y}, v_{0z}, v_{1x}, v_{1y}, v_{1z})$ and $\mathbf{x}_{s1} = (v_{1x}, v_{1y}, v_{1z}, v_{2x}, v_{2y}, v_{2z})$ and consider derivatives with respect to the stretching energy. Note that the central vertex \mathbf{v}_1 receives a force contribution from both stretching elements. By iterating over all force elements and accumulating their force contributions to their degrees of freedom, we obtain the net forces used to step the system.

The derivatives above are non-trivial to derive, so we simply use the results provided in Bergou et al. [1]. Therein we are given gradients of the twist and curvature with respect

to edges. Defining $\tilde{\mathbf{t}} = (\mathbf{t}^{i-1} + \mathbf{t}^i)/\chi$, $\tilde{\mathbf{d}}_\alpha = (\mathbf{d}_\alpha^{i-1} + \mathbf{d}_\alpha^i)/\chi$, with $\chi = 1 + \mathbf{t}^{i-1} \cdot \mathbf{t}^i$, we have

$$\begin{aligned}\frac{\partial m}{\partial \mathbf{e}^{i-1}} &= \frac{\kappa \mathbf{b}}{2|\mathbf{e}^{i-1}|}, \\ \frac{\partial m}{\partial \mathbf{e}^i} &= \frac{\kappa \mathbf{b}}{2|\mathbf{e}^i|}, \\ \frac{\partial \kappa_1}{\partial \mathbf{e}^{i-1}} &= \frac{1}{|\mathbf{e}^{i-1}|}(-\kappa_1 \tilde{\mathbf{t}} + \mathbf{t}^i \times \tilde{\mathbf{d}}_2), \\ \frac{\partial \kappa_1}{\partial \mathbf{e}^i} &= \frac{1}{|\mathbf{e}^i|}(-\kappa_1 \tilde{\mathbf{t}} - \mathbf{t}^{i-1} \times \tilde{\mathbf{d}}_2).\end{aligned}$$

The edge partials of κ_2 are obtained by making the substitutions $\kappa_1 \rightarrow \kappa_2$, $\tilde{\mathbf{d}}_2 \rightarrow -\tilde{\mathbf{d}}_1$. Starting with these expressions, a judicious use of chain and product rules yield the vertex and angle partials of the energies.

2.8 Damping Model

Realistic sounds are sensitive to the kind of damping model used. The most commonly used model in the context of physics-based audio synthesis is the *Rayleigh damping* model, which assumes a damping force \mathbf{F}_d proportional to the velocity, namely,

$$\mathbf{F}_d = -\mathbf{C}\mathbf{v}.$$

The damping matrix $\mathbf{C} = \alpha\mathbf{M} + \beta\mathbf{K}$ is taken to be a linear combination of the mass and stiffness (force Jacobian) matrices, where the parameters α and β are tuned by hand to get a desired sound. Due to the complexity of the bending force Jacobian, we forego its computation and consider the mass matrix alone. A value of α between 10^{-5} s^{-1} and 10^{-4} s^{-1} was found suitable for nicely decaying sounds.

2.9 Validating the Force Expressions

The force expressions obtained above are rather non-trivial, so a reliable, systematic method of testing them is desired. Given the many mathematical entities to compute (frames, curvatures, twists) and the large number of possible configuration of even a minimal system, we forego a traditional unit test approach and, inspired by the definition of force as the gradient of an energy, use numerical finite difference approximations of the forces to validate the analytic expressions. We verify that the finite difference approximations converge to the analytic expressions for a number of procedurally-generated simple test systems, confident that the energy expressions are simple enough to test by hand.

2.10 Simulation Loop

We use a symplectic Euler scheme to integrate the equations of motion. Given accelerations, velocities and positions \mathbf{a}_n , \mathbf{x}_n , \mathbf{v}_n at the n 'th timestep t_n , the state at the next timestep $t_{n+1} = t_n + \Delta t$ is given by

$$\begin{aligned}\mathbf{v}_{n+1} &= \mathbf{v}_n + \mathbf{a}_n \Delta t, \\ \mathbf{x}_{n+1} &= \mathbf{x}_n + \mathbf{v}_{n+1} \Delta t.\end{aligned}$$

This scheme has reasonable stability properties for stiff systems and is very easy to implement compared to higher-order explicit schemes or implicit schemes. While implicit schemes are preferred for very stiff systems, of which realistic elastic rods are an example, the numerical damping of high frequencies leaves them unsuitable for audio synthesis applications. Indeed, Schweickart et al. [19], while using an implicit Newmark scheme, step their simulation at very high frequencies (80-120 kHz) anyway. Thus, we avoid the difficulty of implementing implicit integrators and stick with a symplectic scheme with sufficiently high simulation frequency.

Once the vertex positions and edge angles are stepped, we update the reference frames by time-parallel transport. The new material frames are then obtained by rotating the new reference frames about the tangent by the new edge angle.

Chapter 3

Audio Synthesis

Synthesizing the sound produced by the motion of a body with volume Ω amounts to solving the wave equation for the acoustic pressure field $p(r, t)$ outside the body,

$$\frac{\partial^2 p}{\partial t^2} = c^2 \nabla^2 p(r, t), \quad r \in \mathbb{R}^3 \setminus \Omega,$$

supplemented by appropriate boundary conditions determined by the motion of the surface $\partial\Omega$ of the body and behaviour at infinity. We take the speed of sound in air $c = 340$ m/s. A full solution to this problem, involving the detailed modelling of the surface motion of a rigid body, its coupling to the air, and subsequent propagation of the acoustic field to the listener was implemented by O'Brien et al. [15], but is computationally expensive and overkill for our purposes. We seek a simpler model that captures the salient features of the bending modes associated with the motion of the rod's centerline, ignoring the contribution of the stretching and twisting modes. We follow the discussion presented by Schweickart et al. [19], based on material in Howe [7].

3.1 Sound Radiation Model

We approximate the far-field acoustic pressure due to the motion of a compact, rigidly vibrating rod element as a dipole source. We make a few assumptions to arrive at a particularly simple expression for the acoustic pressure of said source:

- Acoustic compactness - the rod diameter d is much smaller than the wavelengths of interest. For example, requiring a upper bound of 10 kHz on the frequency for $c = 340$ m/s requires the rod diameter $d \ll 3.4$ cm. The cross-sections of our systems are on the order of 10 mm.
- Far-field pressure - we require the distance to the listener be larger than the wavelengths of interest. For example, taking the lowest frequency to be 100 Hz requires a minimum listener distance of 3.4 m, which is reasonable for computer graphics applications.
- Cross-section symmetry - we only consider centerline cross-sections with $x - y - z$ symmetries. Indeed, all of our test systems assume elliptical cross-sections.
- Slowly-varying cross section - the dipole model ignores the vibration of the rod element parallel to its centerline. All of our test systems assume a uniform cross-section.

Under these assumptions, the contribution to the acoustic pressure from the i 'th rod element to a listener at a position \mathbf{r}_i relative to the element, at time t , is given by

$$p_i \left(t + \frac{r_i(t)}{c} \right) = \frac{\rho}{4\pi c r_i(t)} \hat{\mathbf{r}}_i^T \mathbf{R}_i(t) \mathbf{D}_i \mathbf{R}_i(t)^T \ddot{\mathbf{x}}_i.$$

We take the density of air $\rho = 1.225$ kg/m³. The jerk $\ddot{\mathbf{x}}_i$ is computed using a forward-difference stencil. The matrix \mathbf{R}_i approximates the rod's material frame at the i 'th vertex, and is constructed as follows: Set $\mathbf{d}_{i,3} = (\mathbf{d}_3^{i-1} + \mathbf{d}_3^i)/|\mathbf{d}_3^{i-1} + \mathbf{d}_3^i|$, space-parallel transport $\mathbf{d}_1^i, \mathbf{d}_2^i$ to $\mathbf{d}_{i,3}$ and rotate by $(\theta^{i-1} + \theta^i + \underline{m}_i)/2$ to obtain $\mathbf{d}_{1,i}, \mathbf{d}_{2,i}$. Finally, the dipole matrix

D captures the effect of the cross-section geometry on the dipole source far-field. It is computed on a per-element basis and, under the assumptions above, is given by

$$\mathbf{D} = \oint_S (\mathbf{y} - \phi^*(\mathbf{y})) \mathbf{n}(\mathbf{y})^T dS_{\mathbf{y}}.$$

The integral is taken over the surface $S = d\Omega$ of the rod element, with $\mathbf{y} \in S$ and surface normal $\mathbf{n}(\mathbf{y})$. The velocity potential $\phi^*(\mathbf{y}) \in \mathbb{R}^3$ is a time-independent function that, component-wise, satisfies

$$\begin{aligned} \nabla^2 \phi^*(\mathbf{r}) &= 0, \quad \mathbf{r} \in \mathbb{R}^3 \setminus \Omega, \\ \phi^*(\mathbf{r}) &\rightarrow 0, \quad |\mathbf{r}| \rightarrow \infty, \\ \frac{\partial \phi^*(\mathbf{r})}{\partial \mathbf{n}} &= \mathbf{n}, \quad \mathbf{r} \in S. \end{aligned}$$

Noting that a slowly-varying cross-section implies $\mathbf{n}(\mathbf{y}) \perp \hat{\mathbf{z}}$, we get that $\mathbf{D}\hat{\mathbf{z}} = 0$, so we only need to compute the 2×2 top-left submatrix of \mathbf{D} corresponding to the dipole source generated by $x - y$ motion. For a rod element of length Δz , with 2D cross-sectional boundary Γ , the elements of this submatrix are given by

$$D_{ij} = \Delta z \oint_{\Gamma} (y_i - \phi_i^*(\mathbf{y})) n_j(\mathbf{y}) d\Gamma_y, \quad i, j \in 0, 1.$$

Thus, computation of \mathbf{D} reduces to solving a pair of 2D exterior Laplace problems with Neumann boundary conditions to determine the values of ϕ^* on Γ .

3.2 Dipole Source Implementation Details

While a boundary element method approach is appropriate for arbitrary Γ , for the purposes of illustration we consider a simpler finite-difference approach, assuming that the cross section S is sufficiently well-approximated by a collection of rectangles. The exterior Laplace problem may then be recast as a constrained optimization problem, namely




Figure 3.1: Solution to the exterior Laplace problem for different cross-sections. The far-field is insensitive to the specific cross-section.

to minimize $f(\mathbf{x}) = \mathbf{x}^T \nabla^2 \mathbf{x}$, with constraints defined by the boundary conditions. Then the optimal solution \mathbf{x}^* with corresponding Lagrange multipliers λ^* is obtained as the solution to

$$\begin{pmatrix} \nabla^2 & \mathbf{C}^T \\ \mathbf{C} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{x}^* \\ \lambda^* \end{pmatrix} = \begin{pmatrix} 0 \\ \mathbf{b} \end{pmatrix},$$

where ∇^2 is understood to be the matrix of the discrete Laplacian and the constraints are represented by the linear system $\mathbf{C}\mathbf{x} = \mathbf{b}$. Figure 3.1 shows velocity fields (one component) for different cross-sections.

As the elastic force stiffnesses of our rod dynamics are derived assuming elliptical cross-sections, we restrict our acoustic radiation model likewise. For a rod with elliptical cross-section with axes a, b aligned along the z -axis, the matrix is given by

$$\mathbf{D} = 2\pi\Delta z \begin{pmatrix} a^2 & 0 & 0 \\ 0 & b^2 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

We write these pressure values to a buffer which is then serialized for off-line rendering. In general, $t + r_i(t)/c$ will not be an integer value, and so we use a linearly-interpolated buffer write for non-integer propagation times (Figure 3.2).

Recording the pressure at the listener location is the last step in the audio synthesis pipeline and may be made more sophisticated by some spatialization scheme, such as stereo capture, a more sophisticated microphone array, or a numerical head-related trans-




Figure 3.2: Interpolated buffer write: the index corresponds to a given timestep in the simulation, while the offset accounts for the propagation delay. In general the offset will not be an integer value, so we write the pressure sample at time $t + r_i(t)/c$ to buffer locations at $\text{index} + \text{floor}(\text{offset})$ and $\text{index} + \text{floor}(\text{offset}) + 1$ with weights $1 - \alpha$, α , respectively, where $\alpha = \text{offset} - \text{floor}(\text{offset})$.

fer function (HRTF) that models the filtering of sound by the head, ears, and shoulders. We keep things simple and record a stereo buffer, capturing the interaural time difference for minimal spatialization. The stereo channels are recorded using two mono listeners and are then mixed down for offline rendering of a raw audio file.

Chapter 4

Results

Having laid the foundations of elastic rod dynamics and dipole-based audio synthesis from said dynamics in the previous chapters, we now apply these concepts to simulate some non-trivial test systems.

4.1 Ribbons, Polymers, and Slinkies

We first validate the force model with a few systems that capture its salient features in isolation before combining them together in a physically realistic, non-trivial model of a slinky.

Twisted ribbon: A simple system that cleanly illustrates the relation between reference and material frames is a “ribbon” that is uniformly twisted along its centerline (Figure 4.1a). We construct a straight centerline with an exaggerated rectangular cross-section, take the reference frames to be the euclidean basis vectors, and set the edge angles such that the last edge of the ribbon is twisted 180° relative to the first edge, which is clamped. Releasing the ribbon from this rest configuration causes twist waves to propagate along the centerline. Additionally, one can consider twisted equilibrium configurations which are then further twisted out of equilibrium to ensure this case is handled correctly. We




Figure 4.1: Reference and material frames of three systems. a) Twisted ribbon with material frames (CMY) twisted relative to reference frames (RGB). b) Circular slinky with reference frames generated by space-parallel transport. c) Flat slinky with reference frames set to match discretized Frenet frame.

did not choose physical parameters for the mass and stiffnesses as this system was for illustrative purposes only.

Polymer: Due to its initial configuration, the dynamics of the twisted ribbon are governed by the twisting energy alone. In contrast, we consider a bent “polymer” whose dynamics are governed by the stretching and bending energies, and whose undeformed equilibrium configuration has some non-zero curvature (Figure 4.2). By pinning one end and increasing the stiffness, we can approximately model a rigid rod, which is demonstrated nicely by turning on gravity and allowing the rod to swing like a pendulum. Again, we did not consider physically realistic parameters for this demo.

Cable: As a special case of the polymer, we consider a “cable” with zero equilibrium curvature. Subject to gravity, the cable comes to rest in the familiar shape of a catenary (Figure 4.3) .

Circular slinky: Moving towards a physically-realistic system whose dynamics are governed by all three energies, we consider a slinky with circular cross-section (Figure 4.1b). The centerline is given by a helix, and the initial reference frames are constructed by space-parallel transport of a conveniently chosen reference frame on the first edge. We take the initial material frames to coincide with the reference frames, so that there is no




Figure 4.2: A “polymer” with non-zero equilibrium curvature.




Figure 4.3: A cable under gravity comes to rest in the shape of a catenary.

equilibrium twist. Note that these frames do *not* correspond to the Frenet frame of a helix; the discrete analogue of this will be considered in the next example. We clamp the first edge of the slinky and let it fall from rest under gravity. Happily, we reproduce the characteristic wobble of a real slinky under similar conditions.

Flat slinky: Finally, we consider a slinky with a flat cross-section. To model the equilibrium configuration appropriately, we take the initial material frame to approximate the Frenet frame of a helix (Figure 4.1c). To this the end, we take $\mathbf{d}_3 = \mathbf{t}$, and take \mathbf{d}_1 to be the normal component of the vector parallel to the centerline of the helix. Then $\mathbf{d}_2 = \mathbf{t} \times \mathbf{d}_1$. This equilibrium configuration, in contrast to the circular slinky, has a non-zero unde-

formed twist, which leads to subtly different dynamics. These two examples highlight the importance of choosing the material frames correctly; setting the material frames incorrectly can lead to behaviour wildly different from what might be expected, even if the centerline is correct.

4.2 Rubber Bands and Metal Strings

The above examples do not lend themselves well to the kind of controlled activation required for testing the audio synthesis model. We would like to have some control over the initial displacement from equilibrium, the tension, and the location of any “pluck” force. To this end, we consider a collection of physically realistic rubber bands and metal strings. Since these two materials have very different physical properties, a single physical configuration can give rise to wildly different sounds. Thus a small set of parameters can give rise to a rich collection of sounds, in principle.

Initial conditions: We consider two kinds of initial conditions for driving the audio synthesis:

- Initial displacement conditions: the centerline of the rod is initialized to some displaced shape (Figure 4.4). We introduce tension in the rod by stretching the initial configuration away from equilibrium before displacing the centerline perpendicular to its cross section. We consider triangular displacements of varying height and center, and a global cosine profile.
- Activation force: alternatively, we can activate a rod under tension by applying a force to the vertices, modelling a mallet strike. We can model the “softness” of the mallet by taking the force amplitude to be a cosine lobe in time, with variable amplitude and width (Figure 4.5). Similarly, the size of the mallet can be modelled by modulating the force amplitude with a spatial profile centered on a particular vertex.




Figure 4.4: Displacement initial conditions for the rubber band examples. The red and blue spheres mark the locations of the monophonic listeners. a) Large amplitude cosine lobe. b) Symmetric small amplitude piece-wise linear displacement. c) Offset small amplitude piece-wise linear displacement.

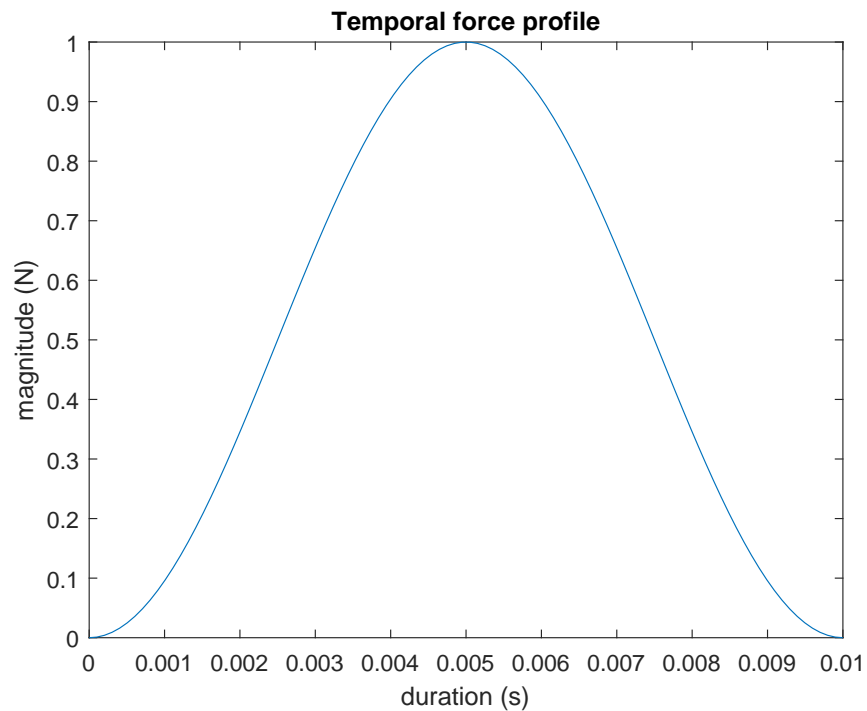



Figure 4.5: Cosine lobe temporal force profile. The duration of the force determines the softness of the virtual mallet.




(a) Rubber band stretched to 40 cm from rest length of 18.6 cm. **(b)** Rubber band stretched to 50 cm from rest length of 18.6 cm.


Figure 4.6: Spectrograms of sound generated from rubber bands of variable tension. In each case the rubber band was initially displaced according to a cosine profile. Though it's not clear from the spectrograms, it can be heard clearly that the rubber band under greater tension produces sound with higher frequency.

Rubber band: We present a number of spectrograms corresponding to a stretched rubber band with variable linear mass density and tension, subject to different initial conditions. Each trial corresponds to 1 second of audio captured at a sampling rate of 20 kHz, which was sufficiently high to capture the frequencies of interest. The rest length of the rubber band is 18.6 cm. The cross-section is elliptical, with major and minor axes 1.6 mm and 0.8 mm, respectively. The reader is encouraged to listen to the audio files included in the supplemental materials to better appreciate the comparisons made herein. In general, changing physical parameters primarily changes the pitch of the sound, whereas the initial conditions and activation forces lead to marked timbral changes in an expected way.

Metal string: The incredible stiffness of real metal strings (Young's modulus on the order of hundreds of GPa) lends them to illustrating the nonlinear effects that are captured by our model, in the form of a time-varying spectrum corresponding to audible pitch shifts. Unfortunately, we found it very difficult to efficiently simulate, in a controlled




(a) Rubber band with linear mass density 900 kg/m^3 .




(b) Rubber band with linear mass density 1120 kg/m^3 .

Figure 4.7: Spectrograms of sound generated from rubber bands of variable linear mass density. In each case the rubber band was initially displaced according to a cosine profile. The rubber band with lower linear mass density sounds with a higher frequency, as expected.




(a) Rubber band released from rest with a centered piece-wise linear profile.



(b) Rubber band released from rest with an offset piece-wise linear profile.

Figure 4.8: Spectrograms of sound generated from rubber bands with different initial displacement conditions. The rubber band with an offset piecewise-linear profile has a richer spectrum, with more energy at higher frequencies.




(a) Rubber band driven by cosine lobe force with unit amplitude and duration 1×10^{-2} s. (b) Rubber band driven by cosine lobe force with unit amplitude and duration 1×10^{-4} s.

Figure 4.9: Spectrograms of sound generated from rubber bands with different activation forces. Both activations were applied to the central vertex, but their temporal duration differs. The “harder” mallet produces sound with an incredibly rich spectrum and a sharp transient at the onset of the sound.


way, real strings that might be found on an instrument like a guitar, and consider a fake material with Young’s modulus on the order of GPa instead. The nonlinear effects are still audible, but the sound produced does not readily correspond to any familiar material.

For reasons of stability, the system is stepped at 10^6 Hz. The corresponding audio signal is downsampled by a factor of 10 before rendering. We consider cosine lobe and centered piecewise-linear initial displacement conditions with large amplitudes. Observe that while the sound generated by the cosine lobe displacement appears reasonably linear, the large amplitude piece-wise linear displacement produces a spectrum with markedly nonlinear effects. Energy moves from lower to higher frequencies as the sound evolves, which manifests as a pitch shift.

Table 4.1 contains geometric information, material parameters and stepsize for the physically realistic systems we simulated. The stepsize is constrained by the stiffness of the system and the sampling rate at which we capture audio information. A valuable



(a) String with large amplitude cosine lobe initial displacement.



(b) String with large amplitude piece-wise linear initial displacement.

Figure 4.10: Spectrograms of sound generated from metal strings with different initial displacement conditions. Large amplitudes lead to nonlinear effects, manifesting as time-dependent spectrums, with energy moving from lower to higher frequencies.

System	Vertices	E (Pa)	ν	μ (kg/m ³)	Step Size (s)
Circular slinky	200	205×10^7	0.29	1.14×10^{-3}	1×10^{-6}
Flat slinky	200	205×10^9	0.29	1.14×10^{-3}	1×10^{-7}
Rubber band	50	3×10^5	0.49	1.95×10^{-3}	5×10^{-6}
Metal string	100	188×10^7	0.5	1.14×10^{-3}	1×10^{-6}

Table 4.1: For each physically realistic system considered, we list the number of vertices defining the centerline, the Young's modulus E , the Poisson ratio ν , the linear mass density μ , and the frequency at which the simulation is stepped.

extension of this work would be to compute detailed timing information for a greater variety of systems and integration schemes.

Chapter 5

Conclusion

We have implemented a framework for the simulation of discrete elastic rods and synthesis of sound from their motion. Based in the Kirchhoff theory of discrete rods, the framework can accommodate a number of non-trivial systems, defined by their centerline and appropriate material frames. We use a symplectic Euler scheme to integrate the system, but our framework can easily be extended to accommodate any integration scheme. In particular, our simulation would benefit from an implicit integrator to achieve stability at larger time steps.

Sound is generated using a far-field dipole source radiation model. The choice of initial displacement and activation forces allows for the generation of a number of different sounds from a few physical parameters characterizing the material composition of the rod. In particular, this model can capture nonlinear pitch glide effects that a linear model cannot. By mixing down the signals captured at two monophonic listening locations, we can render a stereo audio file offline that can then be synchronized with the video captured from the simulation.

5.1 Future Work

A number of improvements can be made to move towards more realistic sound generation, the most fruitful being the implementation of a full Rayleigh damping model that includes the stiffness matrix of the system. This lends itself nicely to an implicit integration scheme such as the linear implicit Euler, which makes use of the stiffness matrix. Additionally, the current sound generation model only uses the bending modes of the rod, so that, for example, a stiff rod dropped on its end would produce no sound. It would thus be valuable to consider a more complicated sound generation model that takes into account the longitudinal and torsional modes of the rod.

Our framework does not feature any collision detection/resolution. This limits the generation of sound from self-contact, such as the “chattering” produced by a slinky. Coupling the rod simulation with rigid body simulations with a full collision detection system would allow for the simulation of complex scenes with many interactions between non-trivial systems.

One glaring drawback of our simulation framework is the time needed to generate even a modest (1-2 seconds) duration of sound. Even using an implicit integrator, as was considered by Schweickart et al. [19] requires tens of seconds to several hours to generate a couple seconds of sound. One approach to speeding up the simulation is to consider multi-scale and reduced models, as in Cirio et al. [6]. Therein the authors consider splitting the system into a low-frequency regime consisting of a few modes that are simulated in a fully nonlinear way, coupled to a high-frequency regime where many more modes are simulated linearly. The authors claim to achieve a tenfold speed up over previous approaches. Since they consider thin shells, it is likely that a much greater speed up could be achieved by applying similar ideas to thin rods.

Appendix A

Supplementary Materials

A playlist of audio and video files referenced in the thesis may be found at the following link: <http://yt.vu/p/PLr2t1MWKNUZxulNIqyIF1XMuNoR0xP4s2>.

Each example is accompanied with a short description. We encourage the reader to have these examples on hand when reading the results chapter.

Bibliography

- [1] BERGOU, M., AUDOLY, B., VOUGA, E., WARDETZKY, M., AND GRINSUN, E. Discrete viscous threads. *ACM Trans. Graph.* 29, 4 (July 2010), 116:1–116:10.
- [2] BERGOU, M., WARDETZKY, M., ROBINSON, S., AUDOLY, B., AND GRINSUN, E. Discrete elastic rods. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 63:1–63:12.
- [3] CHADWICK, J. N., AN, S. S., AND JAMES, D. L. Harmonic shells: A practical non-linear sound model for near-rigid thin shells. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 119:1–119:10.
- [4] CHADWICK, J. N., AND JAMES, D. L. Animating fire with sound. *ACM Trans. Graph.* 30, 4 (July 2011), 84:1–84:8.
- [5] CIRIO, G., LI, D., GRINSUN, E., OTADUY, M. A., AND ZHENG, C. Crumpling sound synthesis. *ACM Trans. Graph.* 35, 6 (Nov. 2016), 181:1–181:11.
- [6] CIRIO, G., QU, A., DRETTAKIS, G., GRINSUN, E., AND ZHENG, C. Multi-scale simulation of nonlinear thin-shell sound with wave turbulence. *ACM Trans. Graph.* 37, 4 (July 2018), 110:1–110:14.
- [7] HOWE, M. S. *Lectures on the Theory of Vortex-Sound*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 31–111.
- [8] JAMES, D. L., BARBIČ, J., AND PAI, D. K. Precomputed acoustic transfer: Output-sensitive, accurate sound generation for geometrically complex vibration sources. *ACM Trans. Graph.* 25, 3 (July 2006), 987–995.

- [9] JAMES, D. L., AND PAI, D. K. DyRT: Dynamic response textures for real time deformation simulation with graphics hardware. *ACM Trans. Graph.* 21, 3 (July 2002), 582–585.
- [10] JAWED, K., NOVELIA, A., AND O'REILLY, O. M. *A Primer on the Kinematics of Discrete Elastic Rods*, 1 ed. Springer International Publishing, 2018.
- [11] KALDOR, J. M., JAMES, D. L., AND MARSCHNER, S. Efficient yarn-based cloth with adaptive contact linearization. *ACM Trans. Graph.* 29, 4 (July 2010), 105:1–105:10.
- [12] LANG, H., AND ARNOLD, M. Numerical aspects in the dynamic simulation of geometrically exact rods. *Applied Numerical Mathematics* 62, 10 (Oct 2012), 1411–1427.
- [13] LANGLOIS, T. R., ZHENG, C., AND JAMES, D. L. Toward animating water with complex acoustic bubbles. *ACM Trans. Graph.* 35, 4 (July 2016), 95:1–95:13.
- [14] LEJEMBLE, T., FONDEVILLA, A., DURIN, N., BLANC-BEYNE, T., SCHRECK, C., MANTEAUX, P.-L., KRY, P. G., AND CANI, M.-P. Interactive procedural simulation of paper tearing with sound. In *Proceedings of the 8th ACM SIGGRAPH Conference on Motion in Games* (New York, NY, USA, 2015), MIG '15, ACM, pp. 143–149.
- [15] O'BRIEN, J. F., COOK, P. R., AND ESSL, G. Synthesizing sounds from physically based motion. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, ACM, pp. 529–536.
- [16] O'BRIEN, J. F., SHEN, C., AND GATCHALIAN, C. M. Synthesizing sounds from rigid-body simulations. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (New York, NY, USA, 2002), SCA '02, ACM, pp. 175–181.
- [17] PENTLAND, A., AND WILLIAMS, J. Good vibrations: Modal dynamics for graphics and animation. *SIGGRAPH Comput. Graph.* 23, 3 (July 1989), 207–214.

- [18] RAGHUVANSHI, N., AND SNYDER, J. Parametric wave field coding for precomputed sound propagation. *ACM Trans. Graph.* 33, 4 (July 2014), 38:1–38:11.
- [19] SCHWEICKART, E., JAMES, D. L., AND MARSCHNER, S. Animating elastic rods with sound. *ACM Trans. Graph.* 36, 4 (July 2017), 115:1–115:10.
- [20] SPILLMANN, J., AND TESCHNER, M. CoRdE: Cosserat rod elements for the dynamic simulation of one-dimensional elastic objects. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Aire-la-Ville, Switzerland, Switzerland, 2007), SCA '07, Eurographics Association, pp. 63–72.
- [21] SUEDA, S., JONES, G. L., LEVIN, D. I. W., AND PAI, D. K. Large-scale dynamic simulation of highly constrained strands. *ACM Trans. Graph.* 30, 4 (July 2011), 39:1–39:10.
- [22] TAKALA, T., AND HAHN, J. Sound rendering. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 211–220.
- [23] VAN DEN DOEL, K., KRY, P. G., AND PAI, D. K. FoleyAutomatic: Physically-based sound effects for interactive simulation and animation. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, ACM, pp. 537–544.
- [24] ZHENG, C., AND JAMES, D. L. Harmonic fluids. *ACM Trans. Graph.* 28, 3 (July 2009), 37:1–37:12.
- [25] ZHENG, C., AND JAMES, D. L. Rigid-body fracture sound with precomputed sound-banks. *ACM Trans. Graph.* 29, 4 (July 2010), 69:1–69:13.